

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES


Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**



(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 November 2002 (14.11.2002)

PCT

(10) International Publication Number
WO 02/091154 A2

(51) International Patent Classification⁷: **G06F 3/00**

(21) International Application Number: **PCT/IE02/00059**

(22) International Filing Date: **2 May 2002 (02.05.2002)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
S010455 **10 May 2001 (10.05.2001) IE**

(71) Applicant (for all designated States except US): **CHANGINGWORLDS LIMITED** [IE/IE]; Trintech Building, South County Business Park, Leopardstown, Dublin 18 (IE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **SMYTH, Barry, Joseph** [IE/IE]; 34 Ferncarrig Rise, Kilgobbin Road, Sandyford, Dublin 18 (IE). **COTTER, Paul** [IE/IE]; 4

Taney Road, Dundrum, Dublin 14 (IE). **MCKENNA, Elizabeth, Ann** [IE/IE]; "Kilcorn", Mullaghland, Mullagh, County Cavan (IE). **HANBY, Nigel, David** [IE/IE]; 5 Terenure Park, Dublin 6W (IE).


(74) Agents: **COYLE, Philip, A.** et al.; F.R. Kelly & Co., 27 Clyde Road, Ballsbridge, Dublin 4 (IE).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: INTELLIGENT INTERNET WEBSITE WITH HIERARCHICAL MENU



(57) Abstract: An internet website which presents a hierarchical menu structure to users includes a personalisation engine 14 to automatically modify the menu structure for each user at least according to that user's previous browsing history at the site. In one embodiment the engine 14 provides, as menu options for a selected menu page of the hierarchical menu, a plurality of the most probable links which the user may wish to reach from the selected menu page. In another embodiment there are provided, as additional menu options for a particular menu page, a plurality of links which are similar, as determined by meeting a predetermined similarity metric, to menu options present on the particular menu page other than through meeting the similarity metric.

**Declarations under Rule 4.17:**

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW, ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent

(AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)
— of inventorship (Rule 4.17(iv)) for US only

Published:

- without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Intelligent Internet Website with Hierarchical Menu

This invention relates to an internet website, and is especially relevant to an internet website of the kind generally known as a portal or gateway, which is a website that is or proposes to be a starting site for users when they first connect to the internet. Examples of major internet portals are Yahoo, Excite and Netscape.

Internet sites or services almost always present a hierarchical menu structure to users. As is well understood, a hierarchical menu structure is one comprising a top level menu page, usually the home page of the portal, having links (selectable menu options) to a plurality of sub-pages each being either a further menu page or a "link" page (i.e. a page of ultimate interest to the user and containing no further menu options of the portal's menu structure). Each further menu page itself has further menu options to a plurality of further sub-pages each of which may be a yet further menu page or a further link page, and so on to the lowest level of the hierarchy which will all be link pages. It is to be noted that a link page may itself contain menu options, for example, if the link page is the home page of a financial website that may have its own menu structure, but these are not under the control of the operator of the internet site and are not considered part of its menu structure.

A problem with existing portals is that the menu structure is fixed. This means that a user may have to "drill down" through several menus and sub-menus to arrive at a link page of interest, and however often the user accesses that page he will always have to go through

that process. Of course, where possible a user can "bookmark" the relevant page for direct access, but he may not always remember to do that and bookmarking may not be available on some WAP devices. In any event, the
5 bookmarking of individual pages does not override the fixed menu structure, which overall may not coincide with the desired browsing habits of the user.

The problem is exacerbated as we move towards the
10 wireless internet and mobile information access. New devices such as mobile phones are limited in terms of their functionality and screen-sizes and this ultimately limits their usefulness as information access devices. For example a mobile phone can have a screen size that is
15 200 times smaller than a typical PC screen, offering the user a narrow access window over the information space, and making it all the more important that every pixel of screen space is used to carry relevant information to end-users. In short, ensuring that mobile users gain
20 access to the right information at the right time, both quickly and easily, is key to the future success of the mobile/wireless information revolution, including the current generation of WAP-enabled services and the next generation of 3G information services.

25 A traditional WAP portal provides the user with a predefined set of navigation menus to bring the user from the home page to their information goal, be it cinema or TV listings, news and weather, etc. The problem as it
30 exists at the moment is that each user is presented with the same front-end and typically must follow a sequence of scrolls and clicks in order to locate their desired information - a "one size fits all" navigation solution.

It is an object of the invention to overcome or mitigate this problem.

Accordingly, the invention provides an internet website which presents a hierarchical menu structure to users, the site having a default menu structure and means to automatically modify the default menu structure for each user over time according to at least that user's previous browsing history at the site.

10

The invention further provides a method of operating an internet website which presents a hierarchical menu structure to users, the method including automatically modifying a default menu structure for each user according to at least that user's previous browsing history at the site.

15

The invention further provides a computer program for use at an internet website which presents a hierarchical menu structure to users, the method including automatically modifying a default menu structure for each user according to at least that user's previous browsing history at the site.

20

While of primary utility in respect of WAP-enabled portals, the invention is applicable to Internet portals in general.

25

Embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings, in which:

30

Figs 1(a) to 1(e) show the sequence of user actions required to access a particular item of information on a WAP phone according to the prior art;

- 5 Figs 2(a) and 2(b) show the sequence of user actions required to access the same item of information on the WAP phone from a WAP portal embodying the invention;

Fig. 3 is a block diagram of the system architecture of a first embodiment of the invention incorporated in a WAP portal;

Fig. 4 is a screen shot of the graphical user interface (GUI) to the menu manager of Fig. 3;

15

Fig. 5 are screen shots illustrating the use of the menu manager to create and configure new pages for personalisation;

- 20 Fig. 6 are screen shots illustrating the use of the customisation engine of Fig. 3;

Figs. 7 to 10 are diagrams explaining the operation of the personalisation engine for a simple menu structure;

25

Fig. 11 are child-parent and parent-child tables;

Fig. 12 is a block diagram of a second embodiment of the invention; and

30

Figs. 13 to 18 are various tables which illustrate the operation of the second embodiment.

In order to better understand and appreciate the problem addressed by the embodiments of the invention, consider the simple example of the use of a WAP phone shown in Fig. 1. In order to access today's cinema listings in their local cinema (the Ormonde cinema in Stillorgan, Dublin) a user must take the following sequence of actions:

- 10 * Fig. 1(a): scroll and select the [Entertainment] option on the home menu (option 7).
- * Fig. 1(b): scroll and select the [Cinema] option on the entertainment menu (option 4)
- 15 * Fig. 1(c): select the [Cinema Times] option in the cinema menu (option 1)
- * Fig. 1(d): scroll and select [Dublin] (option 9)
- 20 * Fig. 1(e): finally, select [Ormonde] (option 8)

This is indicative of a typical user session on current WAP portals and even though this user may query the Ormonde cinema listings on a regular basis the WAP portal cannot modify its structure to better serve this particular user.

Referring to Fig. 2, given the user's preference for using their WAP phone as a source of cinema listings from the Ormonde, the embodiment now to be described will promote the [Ormonde] menu option to a higher position within the WAP menu structure - for example the top level [Entertainment] option would be promoted from position 7 to within the top 3 positions (that is, the first screen

the user sees), Fig. 2(a), with the [Ormonde] option promoted to the top of the entertainment menu, Fig. 2(b). In the example we also see that the user's preference for movies is reflected in the promoted [TV Movies] option from the TV listings section of the WAP site. Of course, the user can also be given more direct access to other related site sections such as the remaining cinema and TV choice.

- 10 Fig. 3 is a block diagram of the system architecture of the first embodiment of the invention incorporated in a WAP portal. It will be understood that the full architecture of the portal is not shown, only those elements necessary to explain the embodiment. The
- 15 embodiment comprises a number of different software components including a menu manager 10 and a menu server 12 including a personalisation engine 14 and a customisation engine 16.
- 20 The menu manager 10 is a user-facing application designed to allow non-technical staff 18 to rapidly develop, configure, and deploy fully personalised information services, whether WAP- or Web-based. The personalisation engine 14 is a sub-system controlling the generation of
- 25 personalised menu pages and the profiling of user behaviour. The client 20 is typically a standard Web or WAP browser. The customisation engine 16 is a sub-system providing end-users 22 with a client application to allow them to fully customise parts of the portal in order to
- 30 fix or override automatic personalisation. The client 24 is typically a standard Web browser.

Typically, the menu manager 10 is used in the following way. An operator 18 first creates and defines the core

content and default structure for their portal using menu manager's rapid prototyping facilities (this will be described with reference to Fig. 4). The core content is stored in a content database 26. Second, still using the menu manager 10, the operator defines the basic personalisation properties including the number of menu options to be seen at each level and the sensitivity of the portal to personalisation (this will be described with reference to Fig. 5). Third, the final portal is deployed and made available to the end-users 22.

Once the portal has been published users can manually customise the portal structure by interacting with the customisation engine 16 (this will be described with reference to Fig. 6). Very briefly, this allows end-users to pre-select certain menu options of the portal site and to promote these to higher positions in the site menu structure. For example, a movie fan might specify that the menu link to their local cinema should be contained within the top-level menu rather than being hidden in a lower-down menu. In addition, the customisation engine allows users to remove parts of the portal site that are not relevant to their needs, and also allows users to specify that certain regions of the site should not be subject to automatic personalisation. Generally speaking, the customisation engine is made available to the user through a standard Web browser 24 and includes an intuitive graphical user interface to allow them to make these customised changes.

Once the portal has been published the personalisation engine 14 begins to monitor and track each individual user's activity and actively modifies each browsing experience according to that user's browsing history.

This means that each time a user requests a new page through their browser, the menu (navigation) options to be displayed on this page will be calculated in real-time based on default options that have been suggested by the portal operator, and on the suggested options recommended by the personalisation engine 14. In practice, popular menu options (for the user) are promoted to a higher position in the hierarchical menu structure so that these regions of the portal site are more immediately accessible to the user in question.

Considering each of the foregoing components in more detail, the menu manager 10 is a user-facing application with an intuitive GUI interface that allows a non-expert to rapidly design and deploy sophisticated information portals. The user can readily define menu and content pages using drag-and-drop functionality. The menu manager will automatically generate and publish the source files (.wml, .html, .xml) that make up the final portal site and even a large-scale portal can be developed and tested in a matter of hours.

A sample screenshot of the menu manager in use is shown in Fig. 4. The main interface is divided into 3 main sections: site structure 30, page properties 32 and error reporting 34. The site structure section 30 allows the operator to define the basic hierarchical menu structure of the portal site using an intuitive tree-based representation and a drag-and-drop style interface, which allows individual pages can be added or deleted or moved to different areas of the site.

Fig. 4 illustrates the construction of one particular portal site rooted at a home page called "e-merge home".

The operator has defined a number of different pages including 1st and 2nd level links. For example the "e-merge home" home page contains a link to a menu page titled, "E-mail + Org", which in turns contains links to
5 information pages called "Messenger", "WAP Mail", etc.

As the operator defines the basic site structure he can also specify the detailed properties of each page in section 32, including:

10

* Title name: The page title as it appears in the menu manager.

15

* Screen label: The page title that will appear on the client browser once the portal has been published, if it is different to the title name. Very often title names will be inappropriate for use as screen labels especially if they are derived from some internal page coding scheme as is usually the case.

20

* Promotion name: The operator can specify a separate promotion name for a page, which will be used if the page in question has been promoted out of its default menu position. This allows the operator to provide
25 additional context information for a page, which would otherwise be lost during promotion. For example, consider a page from the TV section of the portal with "Top Pics" as its screen label. The page is designed to list the top TV programmes for a given day, but this is
30 only clear as long as the page remains as a link from the TV page. If this "Top Pics" page is popular with a user then it will be promoted to a higher level and it's meaning will no longer be clear. To solve this

problem the operator can specify a promotion name such as "TV: Top Pics".

5 * Body: This property allows the operator to specify a presentation template for the current page in order to control how the page is presented to the end-user.

* URL: The URL of the page in question.

10 * Advert: The operator can include a link to an advertising insert which will be included with the page (for Web content) or included before the page (WAP content). A particular benefit of this is that it allows advertising messages to be targeted in the sense
15 that since the operator can associated specific advertisements with relevant menus, these advertisements will be seen more frequently by interested users since these users will have these menus promoted in their personalised menu. For
20 example, RyanAir may purchase advertising space in the travel section of a portal by positioning their advertising banner at the top level travel menu, which typically leads to a whole host of travel related services. The point is that RyanAir want to target
25 people interested in travel since these are more likely to purchase their flights. Ordinarily the RyanAir advertisement and the Travel menu would be statically placed within the site. However, now that the site is personalised, users with an interest in travel will be
30 see the Travel section promoted to a higher place in the menu hierarchy and are more likely to enter the travel part of the site and see the RyanAir advertisement. The ratio of relevant to irrelevant users that view this ad will increases dramatically.

In addition the operator can specify certain key personalisation properties for each page in order to control when and how the page in question is personalised. For example, Fig. 5(a) illustrates the creation of a new menu page, Fig. 5(b) the creation of a new link page, and Fig. 5(c) the setting of different personalisation properties for these pages. The "new page" windows shown in Figs. 5(a) and 5(b) are opened by clicking on appropriate buttons in the menu manager window of Fig. 4, while the "properties" window of Fig. 5(c) is opened by clicking the "Advanced" button in Fig. 5(a) or 5(b). Specifically the following personalisation properties may be set:

* Max Links: The maximum number of links that may appear on the menu page concerned (this is obviously not applicable to a link page). This tells the personaliser how many personalised suggestions may be made for a given page. For example, if the page in question has max links set to 8, but already has 5 links specified by the operator, the personaliser can only recommend 3 additional links. Of course, if any of the 5 operator links have themselves been promoted then additional personalised recommendations may be made. This property is especially important for the development of portals that will be viewed using screen-limited devices such as mobile phones, since it allows the operator to guarantee a maximum page size for each page.

* Hit Threshold: The hit threshold is a number and is used to indicate the minimum number of hits that a page must receive before it is considered for promotion. The benefit here is that it allows the operator to limit

spurious promotion/demotion during the early stages of user activity, when sufficient statistics have not been collected to ensure reliable personalisation. Typically the hit threshold is between 0 and 10 (0 indicates that the page is always available for promotion/demotion).

* No Promote: The operator can specify that the page in question should not be promoted as part of personalisation by setting the no promote property.

10

* Copy Promote vs Move Promote - If copy promote is set for a given page, then when this page is promoted it does not disappear from its original menu position. For instance, in the earlier example we considered the promotion of the "Top Pics" page from the TV menu. If copy promote is set for "Top Pics" then, after promotion, the TV menu still contains a link to "TV Pics". If move promote is set, on the other hand, then in that example "Top Pics" would no longer be a link from the TV menu after promotion.

20

* Promotion: The promotion property allows the operator to specify limits on the distance a page may be promoted. For example, setting promotion to 1 means that the page in question may only be promoted to the same level as its parent menu page. Setting promotion to 4 means that the page may be promoted by up to 4 levels if appropriate.

25

In general then the menu manager 10 allows the operator to rapidly develop and configure a personalised portal. The application includes full project management facilities allowing the operator to develop and manage multiple portals. It also includes an import facility

30

that allows the operator to import an existing portal structure which has been designed using a different portal building package. This features makes it very easy for an operator to re-deploy existing non-personalised portals as fully personalised portals, without having to re-key the portal content. Finally, once the portal has been designed the operator can compile and publish the final content code.

10 The task of the customisation engine 16 is to provide end-users with the facility to manually customise the portal's menu structure.

The customisation engine is capable of generating a graphical representation of the portal structure for the end-user to manipulate using a standard browser client. The structure is generated to reflect the structure that the user would see while actively using the portal; that is, it takes account of the basic portal structure (operator defined) plus any previous customisations (user specified) plus any current personalised recommendations (as provided by the personalisation engine 14).

The basic customisation operation allows a user to request that a given menu page be moved from its default position within the site to a new position, as indicated by the selection of a new parent menu page. All end-user customisations are stored as part of the user's profile structure in a profile database 28, Fig. 3. Furthermore, all customisations must conform to the configuration settings specified by the operator during the design of the portal. For example, a user cannot move a page into a menu if it violates the maximum number of promotion levels allowed in that page. Also, an operator may

specify that certain pages are non-customisable, in which case these pages may not be repositioned.

An example screenshot of the client is shown in Fig.

5 6(a), which depicts a portion of the main portal site.

Fig. 6(b) shows a pop-up window which is activated when the user selects on a portal page (in this case the "FLC Competition" page). This pop-up window allows the user to change the current location of this page. The current
10 location of the page (its default location) is the "New + Cool" menu. The user may select a new location for this page from the selection list in the pop-up window. For example, selecting "e-merge home" would result in the "FLC Competition" link appearing in the "e-merge home"

15 page. Alternatively the user may want to remove the link from his/her view of the portal, which can be achieved by selecting the "Trash" option from the selection list.

The personalisation engine 14 is the heart of the menu
20 server 12. On the face of it the personalisation engine operates in a similar manner to any normal Web server or WAP server: client requests for pages are received and processed and the requested pages are dispatched.

However, it is also capable of monitoring online user
25 behaviour and of modifying the structure of the portal menu structure to best suit the learned preferences of individual users. That is, the personalisation engine is capable of restructuring the content portal so that relevant pages for these users appear closer to the home
30 page, but without compromising the consistency of the portal itself. In brief the personalisation engine is responsible for the following functions:

- * Tracking and profiling online user behaviour. This means monitoring which pages a user requests (which links he selects) and storing this information in their profile where it can be used as a basis for prediction.
- 5 * Combining online behaviour with user-specified customisations to produce a unified user profile.
- * Adjusting the menu structure of the portal individually for each user on a per-request basis to account for their profiled preferences. That is, when a user requests a given page, the personalisation engine dynamically selects the links for that page by combining information about the default structure of the page with information learned about the user's preferences and their customisations, to select a set of best links for the page.

20 The personalisation engine combines a new and efficient method for storing and tracking user behaviour with a unique algorithm for predicting user preferences in terms of where a user is likely to want to go next based on previous browsing behaviour.

25 The basic profile data structure is the hit table, examples of which are shown in Figs. 7 to 10 which will be described in more detail below. The hit table relates parent nodes (menu pages) to child nodes (menu pages or link pages) and the number of times that each child has been selected by the user. For example, on the right hand side of Fig. 9, we see that menu page A has been accessed 20 times, the link to page B from A has been accessed 17 times, the link from B to E 3 times, and so on.

In fact the embodiment uses two tables: a static hit table that is initialised with respect to the default portal structure, and a user hit table that records the user's particular history on the portal. The former is shown on the left hand side of each of Figs. 7 to 10 and the latter on the right hand side.

The static table is defined by the menu manager 10 to reflect the portal structure as specified by the portal operator/designer. It allows the portal to deliver a standard menu structure early on but this will eventually be over-ridden by the personalised menu once the user's history builds up. The hit values that are set in the static hit table make it possible to control the personalisation latency - low values mean that personalisation kicks in very quickly, while large values make the system less sensitive to user activity. In effect, the static hit table is an assumed access history for a first-time new user. As the user browses the portal his selections are used to update the user hit table which is stored as a part of his profile.

The personalisation engine 14 automatically identifies and provides as menu options for a selected menu page m the k most probable links which the user may wish to reach from that menu page, i.e. the k pages with the highest values for $P(n|m)$ based on the current values in the static and user hit tables, subject to any restraints imposed by the site operator (e.g. max links, hit threshold, no promote, etc.) or the user (e.g. non-customisable page). The notation $P(n|m)$ refers to the probability that link n will ultimately be selected by the user given that they are currently in menu m . This

is done dynamically, i.e. in real time, when the user selects that menu page and is not pre-calculated. Thus it is not necessary for the personalisation engine to calculate an entire menu structure in advance based on the user's previous browsing history, but to only calculate the portions of the menu which the user wishes to visit at the time he selects the menu page concerned.

When a user selects a new menu m the personalisation engine locates the k most probable links from m and uses these to construct m . These k most probable links will be different for different users according to their browsing history. Typically, some of the k links will have been specified as links from m by the portal operator but others will have been promoted from lower levels. In order to calculate the k most probable links, the personalisation engine computes $P(n|m)$ for each of the links that have appeared as descendants of m in the past by using the static and user hit tables.

To compute $P(X1|X2)$ where $X1$ is a child of $X2$ (that is, $X1$ is a link in $X2$) we divide the number of hits $X1$ has received as a child of $X2$ (the sum of its static and user hit table values) by the number of hits $X2$ has received.

Generally, if one wishes to compute $P(X1|Xn)$ where $X1$ and Xn are connected according to the parent-child relationships $(X1, X2, \dots, Xn)$, so that $X2$ is a parent of $X1$, $X3$ is a parent of $X2$, and so on...

Then $P(X1|Xn) = P(X1|X2) \times P(X2|X3) \times \dots \times P(Xn-1|Xn)$

Furthermore, suppose $X1$ also occurs as a child of one of more of $X3, \dots, Xn$, then we compute $P(X1|Xp)$ according to

the same method, where X_p is one of X_3, \dots, X_n which has X_1 as a child and add these probabilities to compute a total probability $P(X_1|X_n)$.

■

- 5 We can compute the probabilities of each child node in the normal way, carrying the probabilities from level to level (as child becomes parent) in the obvious way.

- 10 The key is to know when to terminate this process. In theory it may seem that to correctly identify the k most probable links for a menu m that the algorithm must examine each node that has every appeared as a descendant of m . This will become very expensive in large menu structures. However a more efficient technique is possible, based on the following observation:
- 15

$$P(n|m) \geq P(\text{child}(n)|m) \text{ for all children of } n.$$

- 20 That is, the children of a node n always have lower probabilities relative to a node m than the children of n . Having computed the probabilities associated with the links at a particular level we only need to proceed to the next level via those links that have achieved a probability that is greater than the k th best probability found so far.
- 25

■

- This means that we can find the k most probable nodes for menu m and user u by performing a depth-limited breadth-first search over the menu tree rooted at m . The hit tables provide the tree structure information needed to perform the search. That is:
- 30

$$\text{Children}(m) = \text{StaticHitTable}(m) \cup \text{UserHitTable}(m, u)$$

Where $\text{StaticHitTable}(m)$ means the nodes indexed under m in the static hit table and $\text{UserHitTable}(m,u)$ means the nodes indexed under m in the user hit table for user u .

- 5 The following example describes the personalisation process in detail by illustrating how the personalisation progresses over time in terms of a simple menu structure. We assume a hit threshold value of 3 and that each menu page will be constructed to contain 2 links only.

10

1. New User

Fig. 7 illustrates the situation for a first-time user. The user hit table is empty since no nodes (pages) have
15 been clicked.

Consider the construction of menu page A for the first-time user. The hit tables tell us that there are two direct children, B and C, for node A:

20

$$P(B|A) = (20+0)/(40+0) = 20/40 = 0.5$$

$$P(C|A) = (20+0)/(40+0) = 20/40 = 0.5$$

At this point there is no need to proceed any further
25 because neither B nor C have achieved their hit threshold values of 3 (their user hit table entries are empty indicating a hit value of 0 in each case). Thus links B and C will be used in the menu page A and ordered according to their probabilities (in the case of equal
30 probabilities the ordering is computed with reference to the hit values in the static table).

2. Very Early-Stage User

Now consider the situation depicted in Fig. 8. The same user has been using the portal for a very short time and so has built up some initial user hit values. Consider once again the construction of the menu page A. Computing
5 the probabilities for B and C relative to A we find the following:

$$P(B|A) = (20+3)/(40+5) = 23/45 = 0.511$$

$$P(C|A) = (20+2)/(40+5) = 22/45 = 0.488$$

10

Now at this point B has received 3 user hits, achieving the hit threshold value. In addition, because we are looking to identify two links for each menu page we need only expand nodes whose probability is greater than the
15 probability of the second best node probability so far found. This is because, by definition, none of the children or descendant links of a menu option n in menu m can have higher probabilities than n. In this case the second best node probability so far found is 0.488. Since
20 the probability of B is greater (0.511), and because B has achieved its hit threshold, we need to expand B to compute the probabilities of its children. However, we note that none of B's children have achieved their hit thresholds and so they are not available for promotion.
25 Hence, links B and C once again are presented in menu A, with B appearing before C because of its higher probability.

3. Early-Stage User

30

Fig. 9 depicts the situation for the user, still in the very early stages, but where enough usage information has been gathered to begin to find some personalisation effects.

Consider once again the construction of the menu page A. Computing the probabilities for B and C relative to A we find the following:

5

$$P(B|A) = (20+17)/(40+20) = 37/60 = 0.616$$

$$P(C|A) = (20+3)/(40+20) = 23/60 = 0.383$$

At this point both B and C have achieved their hit
10 thresholds. Since we are looking for 2 links for menu page A we need to expand nodes whose probability exceeds the probability of the second highest probability found so far (that is the probability for C, which is 0.383). Therefore we only need to expand node B (whose
15 probability is 0.616).

Expanding B we compute the probabilities of its children, relative to A:

$$\begin{aligned} 20 \quad P(D|A) &= P(D|B) \times (P(B|A)) = (10+14)/(20+17) \times 0.616 \\ &= 24/37(0.616) = 0.399 \end{aligned}$$

$$\begin{aligned} P(E|A) &= P(E|B) \times (P(B|A)) = (10+3)/(20+17) \times (0.616) \\ &= 13/37(0.616) = 0.216 \end{aligned}$$

25 Now at this point we find that the probability of node D is high enough to make it into the top two nodes overall, overtaking node C. Thus at this stage we can construct menu page A to consist of links to B and D, thus D is promoted to menu A from menu B; again B is presented
30 first with D being second.

It is worth noting at this stage that further rules may dictate which nodes are actually chosen and presented to the user in menu A. For example the network operator may

not want static nodes to be deleted for fear that the user would lose access to these menu options. This is viewed as a post-processing task and does not impact on the operation of the personaliser and hence is not at issue here.

4. Mature User

Finally we look at the situation for a mature user where promotion has taken place over time and been acted on by the user. Fig. 10 illustrates the situation: the user has benefited from the promotion of D to menu A and has clicked on this link 5 times in its promoted position as indicated in the user hit table. Note that in Figs. 7 to 10 the user hit table is simply a record of the user's activity - it does not represent the actual menu structure at the relevant time. For example, in Fig. 10 one might be tempted to interpret the user hit table as indicating that the menu page A has three entries, B, D and C. This is not the case.

Computing menu A for Fig. 10 we get the following probabilities for the children of A (note that now D is a possible child of A):

$$P(B|A) = (20+17)/(40+25) = 37/65 = 0.569$$

$$P(C|A) = (20+3)/(40+25) = 23/65 = 0.353$$

$$P(D|A) = (0+5)/(40+25) = 5/65 = 0.076$$

We expand B further because it is the only node whose probability exceeds the worst acceptable probability so far found (0.353, the probability for C).

Expanding B we compute the probabilities of its children,
relative to A:

$$\begin{aligned} P(D|A) &= P(D|B) \times (P(B|A) + 0.076) \\ &= (10+14)/(20+17) \times (0.569) + 0.076 \\ &= 24/37(0.569) + 0.076 = 0.444 \end{aligned}$$

$$\begin{aligned} P(E|A) &= P(E|B) \times P(B|A) = (10+3)/(20+17) \times (0.569) \\ &= 13/37(0.569) = 0.199 \end{aligned}$$

The important thing to note here is that in computing the
probability for node D we consider its probability
relative to B (its static parent) and its probability
relative to A (its promoted parent), by adding the
appropriate probabilities.

In general, as nodes are promoted and demoted over time
they will generate multiple entries in the user hit table
and in computing probabilities we always take the sum of
the probabilities for a given node over its multiple
ancestral entries - by ancestral we mean that we only sum
probabilities for entries of the node that appear higher
up in the menu structure than our current computation
point. Hence, when we are computing the probability for
D relative to A (at A) we do not consider the second
entry for D relative to B since it is below A. But when
we come to compute the probability of D relative to A (at
B) then we take account of both probabilities.

So now the best two probabilities found so far are 0.589
and 0.444, which correspond to nodes B and D
respectively. Once again these nodes can be chosen for
menu A and ordered according to their probabilities.

Although the foregoing has described only a very simple menu structure, the invention is applicable to menu structures having many levels and page promotion can occur up through many level as their probabilities build up over time. The general algorithm is attached at Appendix A.

As it stands the above algorithm defaults to a copy-promote scheme - promoting node D from menu B to menu A does not mean that the user will not see D in B if they select C.

Consider the example described in Figure 10 where menu A is displayed with links to D and B. Now suppose that the user selects menu B (the original parent of D). In a move-promote scheme, even if D is in the top k nodes from B it should not be displayed because it has already been seen by the user in an ancestral menu (in this case the parent menu, but it could just have easily been a grandparent menu etc.). Under the present scheme D would be repeated as a child of B.

If we do wish to implement a move-promote scheme we need two new tables on a session-by-session basis to track promotions and demotions during that session. The sounds complicated but is actually very simple. One table manages the parent-child relationships (and is a simpler version of the user hit table). The other captures the inverse child-parent relationships. During each session we build these tables according to the nodes that are selected for each personalised menu.

Consider the example in Figure 10 once again, and suppose that the user has just accessed menu A with its two

links, D and B. The user's current child-parent and parent-child tables will contain the information shown in Figure 11.

5 Now, when it comes to building menu B we can test if a given node n has already been seen by the user in an ancestral menu by recursively backtracking up the menu tree, by following the child-parent relationships from the child-parent table and then checking the children at
10 each level from the parent-child table.

For example, first we would follow the parent of B (if there are multiple parents we can arrange it so that the most recent parent is listed first in the parent-child
15 table). The check to see if n is an element of the children of this parent (by using the parent-child table). Given that even exponentially large menu trees have linear depths, this is a very cheap operation especially since a single user session is like to focus
20 on a small part of the menu structure.

Various modifications and refinements of the above embodiment are possible. For example, it is possible to specify that a page is promoted only if its probability
25 exceeds a certain threshold, say 0.4, even though it may be in the top k pages according to the probability calculations.

It is also possible to provide a sensitivity attribute to
30 a node. The sensitivity of a node indicates how sensitive it is to personalisation. A low sensitivity value means that the node is relatively insensitive to personalisation and so will move slowly. The converse is true for a node with a high sensitivity value. This is

controlled in practice by using the sensitivity value as the increment when a given node is selected. For example, if a node has a sensitivity value of 3 then each time it is selected 3 is added to its current total in the appropriate entry in the hit table. One complication with this scheme is that potentially the hit value of a parent node may no longer be the sum of the hit values of its children (if children have sensitivity values different from the parent). For this reason a simple normalisation process is used to ensure that the hit values of child nodes sum to equal the hit values of the parent node. Very briefly the parent's hit values are distributed among the child nodes according to the ratios of actual child hit values. Sensitivity is essentially a way of weighting the probabilities of the nodes.

In summary, the embodiment of the invention comprises an application that allows users to rapidly develop and deploy personalised information services - that is, information services that automatically modify to the needs of individual users based on their online behaviour patterns. The constructed portal is capable of fully modifying to the navigation preferences of individual online users and of ultimately delivering a more efficient and enjoyable mobile experience to these users. The embodiment allows users to benefit from a shorter "click-distance" to their preferred content.

The embodiment provides these benefits by allowing the WAP portal to modify dynamically to the access patterns of an individual user in order to proactively predict that user's likely (and common) information need and so promote the information target to a more accessible position within the WAP portal menu structure.

Referring now to Fig. 12, a second embodiment of the invention retains the personalisation functionality of the first embodiments described above, but adds further
5 function based on what is referred to herein as an affinity recommendation technique. For simplicity, in Fig. 12 the customisation engine 16 is omitted in the view of the menu server 12, as are the menu manager 10, browsing client 20 and customisation client 24. However,
10 all are present and operate as described above. As compared to the first embodiment the embodiment of Fig. 12 includes an affinity recommendation engine 40 which operates as follows.

15 With the affinity recommendations technique, menu options are recognized as being comparable to each other if they both receive a number of clicks from the same set of users. The concept is that if the same set of users are interested in both menu options then there is a high
20 probability that the menu options are possibly related to individual users in similar ways. This differs from existing collaborative filtering techniques in that it incorporates similarity between menu options that are not necessarily conceptually similar. For example, an
25 affinity recommendation technique may identify a high similarity between a technology news menu option and a TV listings menu option even though there is no direct similarity between the content they lead to. This similarity value means that a group of users who view
30 technology news also view TV listings so they have a similar interest in both these topics.

The affinity recommendation engine 40 works on the basis of user navigation patterns. Essentially this navigation occurs across menu options in a mobile portal and the users' navigation patterns are used to cluster menu options together.

The first step in the affinity recommendation technique is to note what users have clicked on each of the menu options. This done by making an initial pass through the user hit tables and building an inverted index keyed by menu options. Each of these menu options has a list of associated profiles that contain these menu options and the number of times they have been clicked by the owners of the profiles. As the number of users increases it becomes infeasible to store this inverted index in memory. Therefore each entry in the inverted index is stored in a separate file. For each menu option a file is created with each line in the file containing the name of a profile and the number of times that menu option has been clicked by the owner of the profile.

For example, the file for the menu option M1 may contain the data shown in Fig. 13, which means User1 clicked M1 10 times, User6 clicked M1 50 times, and so on. In order to reduce the number and size of these files the following threshold values can be used:

Hit threshold - when building up a file for a menu option such as M1 only include those users whose click value for M1 ~~exceeds~~meets or exceeds the hit threshold. The hit threshold is a predefined number which represents the lower limit on the number of clicks a menu option must receive from a user before the user's profile is deemed

useful in computing the affinity values for this menu option.

User threshold - if the number of users added to the file associated with the menu option exceed the user threshold stop adding users to the file. These files may be removed altogether so they don't take part in the affinity recommendation. The reason for removing them is that the menu option is a very common menu so may not provide additional benefits in the affinity recommendation process as it is already high up in most users' personalized menu structure.

When updating the inverted index files, only those hit tables within user profiles that have been accessed since the previous update are examined. If the profiles have not been accessed then the number of clicks associated with hit tables in the profiles will not have changed. As a result the profiles will have no affect on the files that are being updated.

In order to recommend menu options that are similar to other menu options a method for computing the similarity between menu options is needed. This uses the association of profiles to menu options. Consider the similarity between two menu options M1 and M2, Fig. 14. M1 has been clicked by User1, User3, User5, User10 and User11 and M2 has been clicked by User1, User4, User5, User11, and User20. To compute the similarity between M1 and M2 we only need to consider the number of hits given to each menu option by the users that have clicked on both (i.e. User1, User5 and User11).

The similarity between M1 and M2 can be computed using Pearson's Correlation Coefficient, which returns a value between -1 and +1. The higher the value the more similar the menus are and vice versa. The following example
 5 shows the similarity computation between menus M1 and M2 in Fig. 14. The resulting similarity is 0.266 which suggests the menus are not very similar.

Pearson's Correlation Coefficient:

$$r = \frac{\sum XY - \frac{\sum X \sum Y}{N}}{\sqrt{(\sum X^2 - \frac{(\sum X)^2}{N}) (\sum Y^2 - \frac{(\sum Y)^2}{N})}}$$

10

If we consider the example above the similarity is computed as follows. The values of X in the equation are assigned those associated with menu option M1 and the
 15 values of Y are assigned the values associated with menu option M2, Fig. 15.

Taken step by step:

$$20 \quad \sum XY = (10 \times 5) + (20 \times 16) + (15 \times 45) = 1045$$

$$\sum X = 10 + 20 + 15 = 45$$

$$\sum Y = 5 + 16 + 45 = 66$$

$$\sum X^2 = 100 + 400 + 225 = 725$$

$$\sum Y^2 = 25 + 256 + 2025 = 2306$$

$$25 \quad (\sum X)^2 = 2025$$

$$(\sum Y)^2 = 4356$$

$$N = 3$$

$$\text{Similarity} = 0.266$$

When computing the similarity between menu options an "overlap threshold" is needed. This threshold represents the minimum number of users needed to have clicked on both menu options before Pearson's Correlation Coefficient (or any other similarity technique) can be used to measure the affinity between menu options. Any number lower than this threshold would suggest the menu options might not be similar, as very few users have selected both menu options.

10

In the above example Pearson's Correlation Coefficient is used to compute the similarity between the menu options. However, this is only one possible technique for computing a similarity measure. Any well-defined similarity metric can be used to measure the similarity between menu options.

15

The affinity recommendation engine 40 constructs affinity files which are used to maintain the relationship (i.e. the affinity) among menu options. There will be an affinity file for each menu option, which contains a list of other menu options and the relationship between them and the owner of the affinity file. For example, the content of the affinity file for M1 may look like Fig. 16, which means M1 and M2 have a similarity value of 0.266, etc.

20

25

As with the index files, the size of the affinity files may be limited in various ways:

30

1. Partners - only consider the partners of a menu option for inclusion in the affinity files. There are a variety of methods for selecting the partners of a

particular menu option. For example, and referring to Fig. 17, the partners of menu option B can be:

* The descendents (D, E, H, I) of the menu option.

* All nodes in the menu structure, except those (D, E, H and I) that are descendents of the menu option.

* The siblings (c) and the descendents (D, E, H, I) of the menu option.

* Every menu option that is a leaf node (i.e. a link to an external content provider) and a non-descendent of the menu (F, J, K). This will ensure that only leaf nodes from other areas of the menu structure may be recommended along with the menu option. These are recommended because they may be menu options the user has not seen before but may be of interest and because of its similarity to the user's current menu option.

* All other menu options in the menu structure.

* Any combination of nodes in the menu structure.

In the foregoing, the menu structure from which the partners are selected is, of course, the default menu structure for the user, since the personalised menu is only built dynamically as the user clicks through it. However, the affinity recommendation technique is not dependent, and is usable independently of, the personalisation technique described in relation to the first embodiment, and in that case the menu structure from which the partners are defined is just the menu structure of the site.

2. Affinity threshold - only include those menu options that have a similarity value above a certain threshold, i.e. the affinity threshold. The reason for this is to ensure that only menu options that are sufficiently similar to a menu option may be recommended along with that menu option.

10 Affinity recommendations can be combined with any menu structure to provide personalized recommendations. For any menu option there is the potential for recommending any other menu option that is sufficiently (i.e. above the specified affinity threshold) similar. When, as in
15 the present embodiment, the affinity recommendations are combined with the menu options built using the probability model, the simplest approach is to build a personalized menu (based on the probability model) and then include a separate list of personalized affinity
20 recommendations as shown in Fig. 18, where M=menu options based on probability model, A=affinity recommendations.

The site operator can predefine the number of recommendations. To produce a list of personalized
25 affinity recommendations an affinity source is identified. The affinity source is the menu or list of menus that are used in producing the affinity recommendations. An affinity recommendation score is also used which represents the likelihood of a menu
30 option to appear as an affinity recommendation. This score may be simply the affinity value associated with the menu option or may be some combination of probability and affinity values. The following describes a number of

techniques for selecting the affinity recommendations to appear on a selected menu page.

Consider a menu page, M, with three menu options, M1, M2
 5 and M3. The task is to produce a list of k affinity
 recommendations to appear on M. The following describes a
 number of techniques that may be used to select the
 affinity recommendations. Note: Methods A and B use the
 affinity value as the affinity recommendations score.

10

A. Use M as the affinity source.

The simplest approach would be to use M as the source for
 the affinity recommendations. The list of affinity
 15 relations for the menu page M are ordered in decreasing
 order of affinity value and the menu options associated
 with the top k relations are used as the affinity
 recommendations.

20 B. Use M1, M2 and M3 as the affinity source.

With this approach the top k affinity relations for M1,
 M2 and M3 are combined. This combination of relations is
 in turn ordered by affinity value and the top k` are
 25 selected as the affinity recommendations.

$\text{Union}(\text{Affinity}(M1, M2, M3)) = \{A_1^{M1} \dots A_k^{M1}, A_1^{M2} \dots A_k^{M2},$
 $A_1^{M3} \dots A_k^{M3}\}$ where A_1^{M1} is the menu with the highest
 affinity relation to M1 and so on.

30

Affinity Recommendation = the top k menu options from the above set with the highest affinity value.

5 C. Use M1, M2 and M3 as the affinity source but vary the affinity recommendations score.

This method is similar to (B) except it incorporates probability information from the affinity source when selecting the k best affinity recommendations. For each
10 of the top k affinity relations associated with M1, M2 and M3 an affinity recommendations score is computed using the following equation:

$$\text{Score}(A_1^{M1}) = \text{Probability}(M1) * \text{Affinity}(A_1^{M1}, M1)$$

15 Where the Probability(M1) is the probability of M1 appearing on the menu page M (based on the probability model). Affinity(A_1^{M1} , M1) is the affinity relation (i.e. the similarity) between M1 and the affinity recommendation A_1^{M1} .

20

The possible affinity relations are then ordered in decreasing order of affinity score and the top k are used as the affinity recommendations. Note: In the above equation the score is based on the product of the
25 probability and affinity values. Other combinations may be used to compute score.

D. Use various combinations of M1, M2 and M3.

The techniques described in B and C above can be used on various combinations of M1, M2 and M3 such as:

- 5 * Those menu options from M1 to M3 that are part of the default menu page M.
- * Those menu options from M1 to M3 that are not part of the default menu page M but have been promoted to the menu page M using the probability model.
- 10 * Those menu options from M1 to M3 that are not part of the dynamic menu page M but have been customized by a user to the menu page M.

Note: Affinity recommendations selected using method A or using the default menus as the affinity source will be
15 the same for all users. However, with the other techniques described the affinity recommendations appearing on a particular menu page M are personalized, so the menu options selected as recommendations may vary from person to person.

20

The end result, therefore, is a personalised menu 42, Fig. 12, consisting of the personalised menu structure 44 produced by the probability calculation and including in addition the affinity recommendations 46.

25

The invention is not limited to the embodiments described herein which may be modified or varied without departing from the scope of the invention.

Appendix A

Personalize(userSession, parentNode, k)

```

5  {
    expand = true
    staticChildren  $\rightarrow$  the static children of the parentNode
    clickedChildren  $\rightarrow$  nodes clicked by user from this parentNode

10     if (clickedChildren is empty)
        possibleNodes  $\rightarrow$  staticChildren
    else
        possibleNodes  $\rightarrow$  clickedChildren  $\dot{\cup}$  staticChildren

15     possibleNodes  $\rightarrow$  possibleNodes - nodes hidden by user
        possibleNodes  $\rightarrow$  possible Nodes - nodes customised by user to a different level

    if (project implements move promote)
        possibleNodes  $\rightarrow$  possible Nodes - nodes promoted to a different level

20     compute probabilities of possibleNodes
        sort possibleNodes in descending order of probability

    while (expand is true) {
25         expand  $\rightarrow$  false
        kBestNodes  $\rightarrow$  the k nodes in possibleNodes with highest probability
        kthMostProbableNode  $\rightarrow$  kth most probable node in the possibleNodes
        kthProbabilityValue  $\rightarrow$  the probability of the kthMostProbableNode

30         For each of the kBestNodes {

            If (expand is false) {
                currentNode  $\rightarrow$  next node in kBestNodes

35                 if (node has not being expanded before and node is a menu node) {
                    staticChildren  $\rightarrow$  the static children of currentNode
                    clickedChildren  $\rightarrow$  nodes clicked by user from currentNode

                    if (clickedChildren is empty)
40                         possibleNodes  $\rightarrow$  possibleNodes  $\dot{\cup}$  staticChildren
                    else
                        possibleNodes  $\rightarrow$  possibleNodes  $\dot{\cup}$  (clickedChildren  $\dot{\cup}$  staticChildren)

                    if (project implements move promote)
45                         possibleNodes  $\rightarrow$  possibleNodes - nodes promoted to a different level

                    possibleNodes  $\rightarrow$  possibleNodes - nodes hidden by user
                    possibleNodes  $\rightarrow$  possibleNodes - nodes customised by user to a different parent

50                    compute probabilities of possibleNodes
                    sort possibleNodes in descending order of probability
                    expand  $\rightarrow$  true

55                }
            }
        }
    }
}

```

Claims

1. An internet website which presents a hierarchical menu structure to users, the site having a default menu
5 structure and means to automatically modify the default menu structure for each user over time according to at least that user's previous browsing history at the site.

2. An internet website as claimed in claim 1, wherein
10 the modifying means includes means for promoting, as menu options for a selected menu page of the menu structure, a plurality of the most probable links which the user may wish to reach from the selected menu page, subject to any restraints imposed by the site operator or the user.

15 3. An internet website as claimed in claim 2, wherein the probabilities of the links are weighted.

4. An internet website as claimed in claim 2 or 3,
20 further including means for providing user-designated links as menu options for a selected menu page.

5. An internet website as claimed in claim 2, 3 or 4,
25 further including means for limiting the number of menu levels that a link can be promoted within the default menu structure.

6. An internet website as claimed in any one of claims
30 2 to 5, further including means for selecting whether a promoted link is copied or moved to the selected menu page.

7. An internet website as claimed in any preceding claim, wherein the modifying means dynamically calculates

the menu options which appear on a selected menu page when the user selects that menu page.

8. An internet website as claimed in any preceding claim, wherein the modifying means automatically modifies the default menu structure for each user according to both that and other users' previous browsing history at the site.

9. An internet website as claimed in claim 8, wherein the modifying means includes means for providing, as at least one additional menu option for a particular menu page, at least one link which is similar, as determined by a predetermined similarity measurement, to at least one menu option present on the particular menu page other than through the similarity measurement, the similarity measurement being a function of that and other users' browsing history.

10. An internet website as claimed in claim 9, wherein the similarity between two menu options is determined as a function of the number of times each menu option has been selected by users that have selected both.

11. An internet website as claimed in claim 10, wherein a user's selection of a particular menu option is taken into account in the similarity measurement only if the number of times that user has selected that menu option reaches a predetermined threshold.

30

12. An internet website as claimed in claim 10 or 11, wherein a menu option is not considered for provision as an additional menu option if the number of users selecting it exceeds a predetermined threshold.

13. An internet website as claimed in claim 10, 11 or 12, wherein a similarity measurement is made on two menu options only if each menu option has been selected by a
5 predetermined minimum number of users.

14. A method of operating an internet website which presents a hierarchical menu structure to users, the method including automatically modifying a default menu
10 structure for each user over time according to at least that user's previous browsing history at the site.

15. A method as claimed in claim 14, including automatically modifying the menu structure hierarchy for
15 each user according to both that and other users' previous browsing history at the site.

16. A computer program for use at an internet website which presents a hierarchical menu structure to users,
20 the program having code for automatically modifying a default menu structure for each user over time according to at least that user's previous browsing history at the site.

25 17. A computer program as claimed in claim 16, including code for automatically modifying the menu structure hierarchy for each user according to both that and other users' previous browsing history at the site.




FIG. 1




FIG. 2

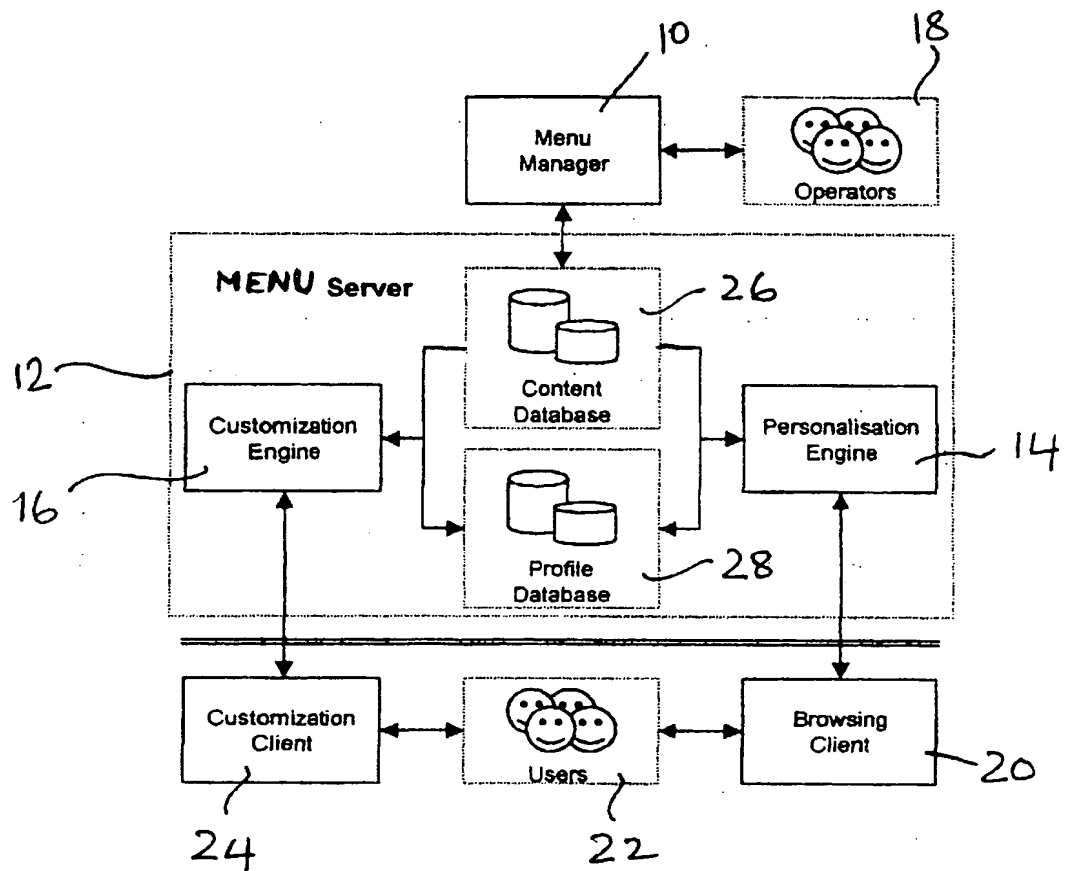



FIG. 3



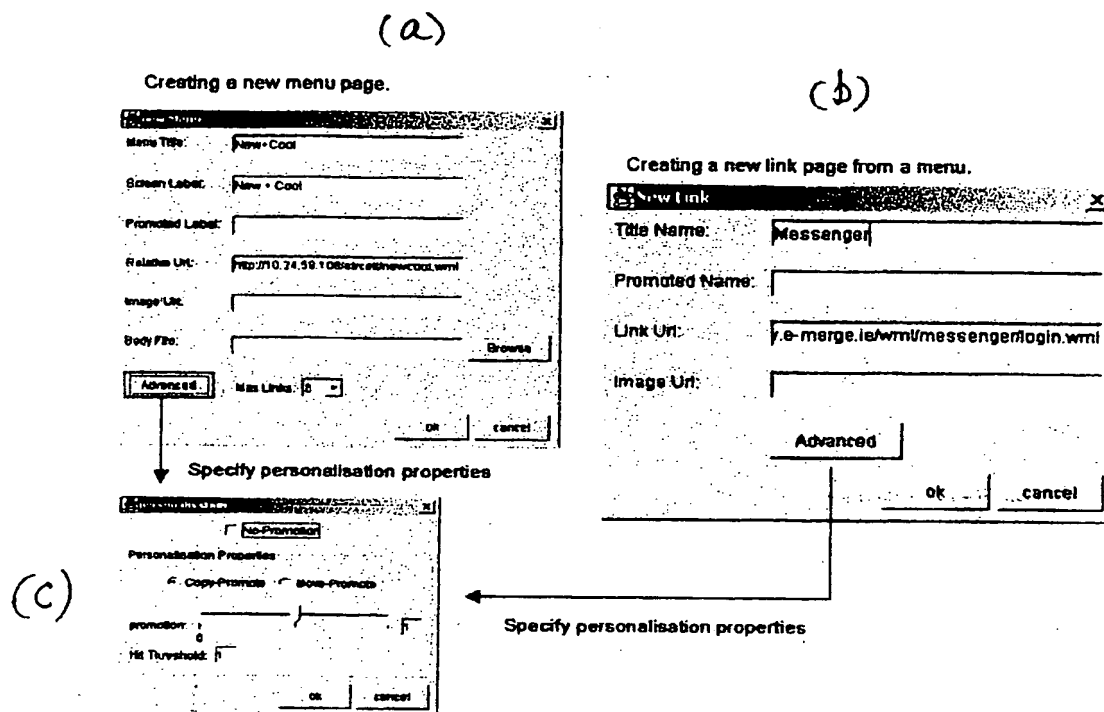


FIG. 5

Parent-child Table				Child-parent Table	
A	B	D	C	B	A
				D	A
				C	A

FIG. 11

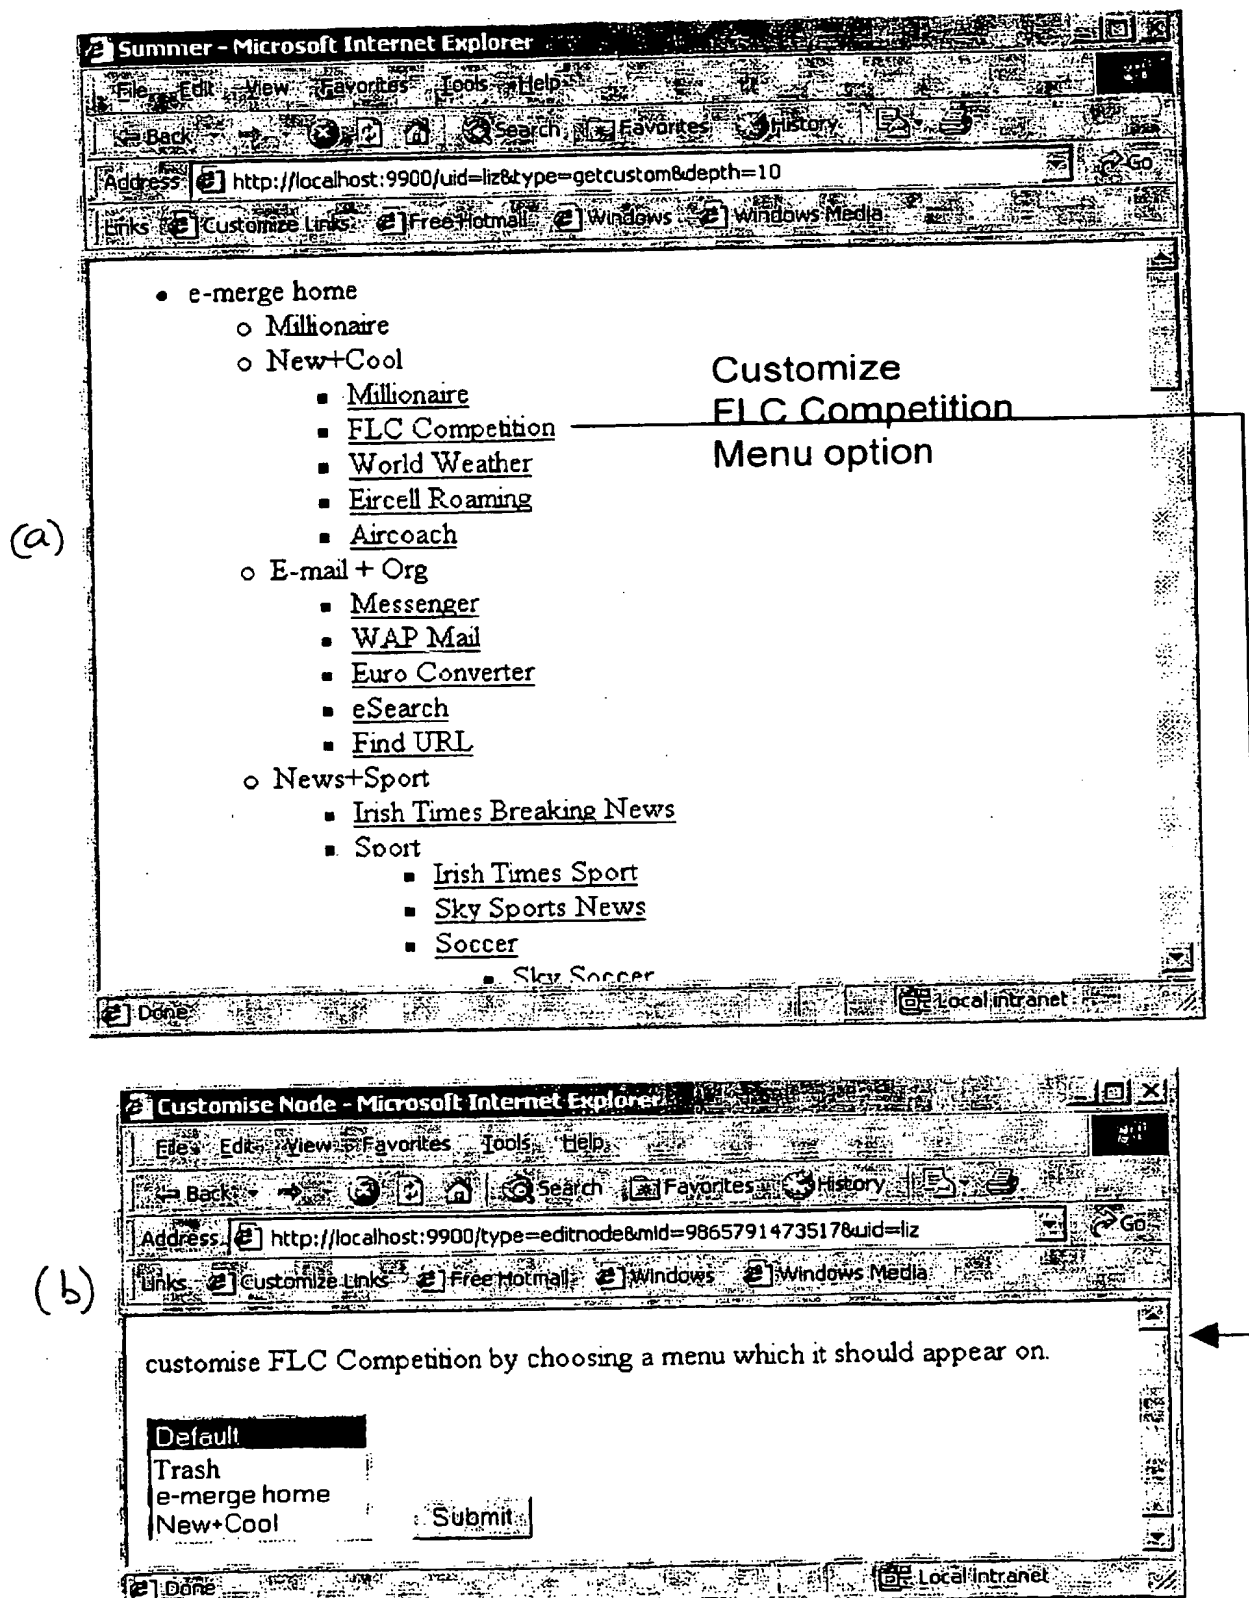


FIG. 6

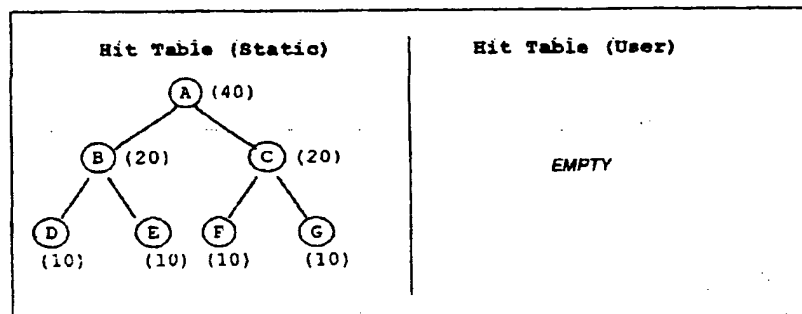


FIG. 7

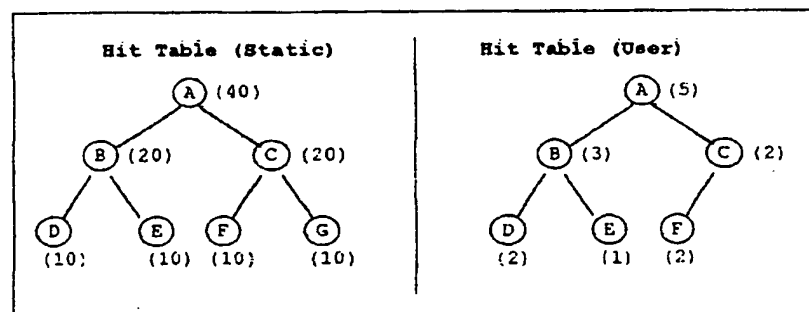


FIG. 8

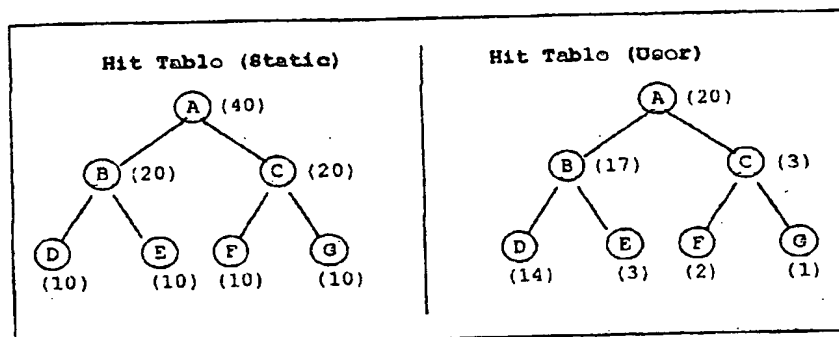


FIG. 9

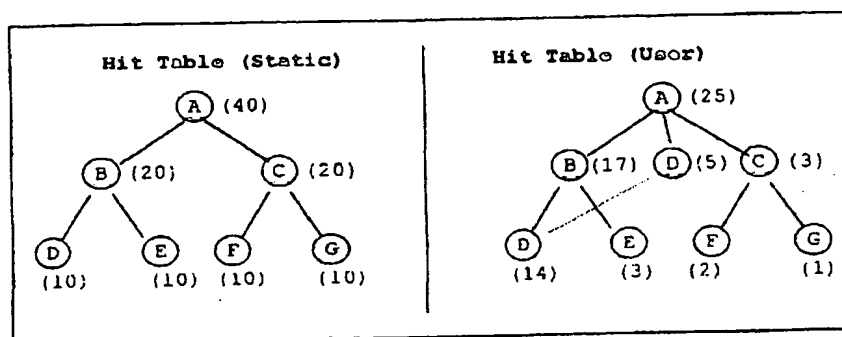


FIG. 10

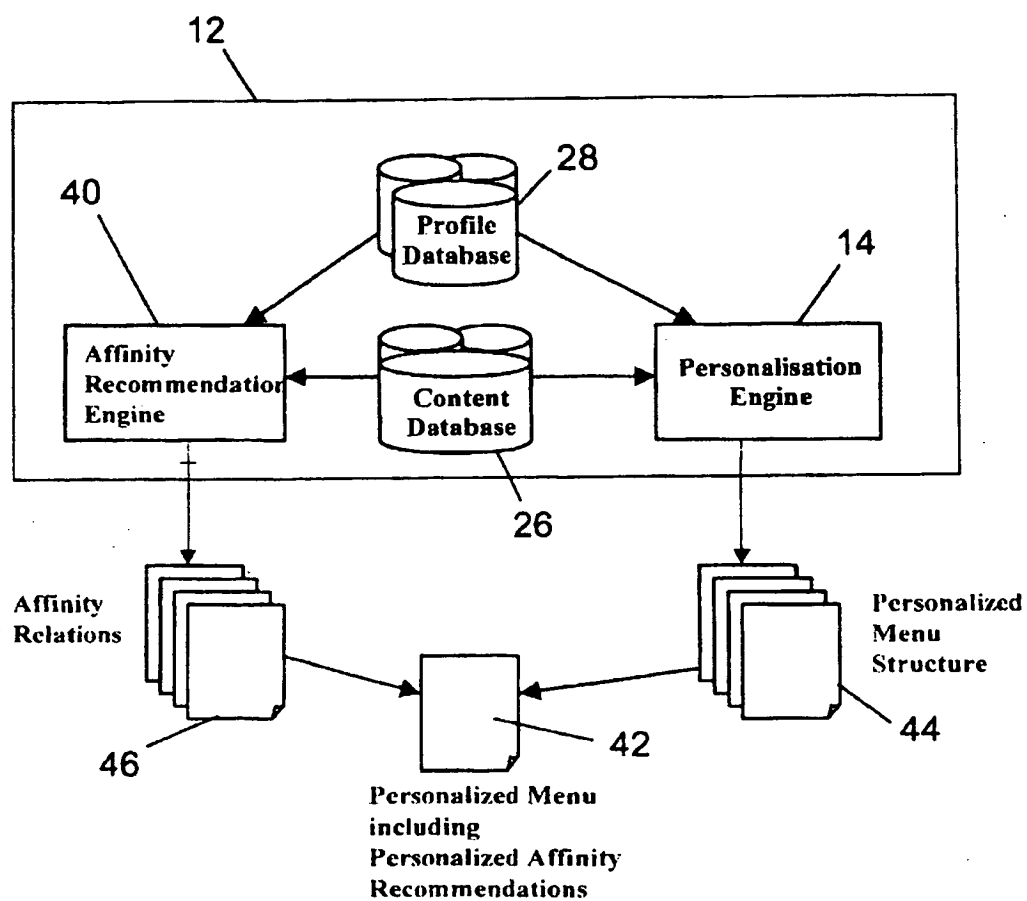


Fig. 12

M1

Users	Hits
User1	10
User6	50
User8	20

Fig. 13

M1		M2	
User	Hits	User	Hits
User1	10	User1	5
User5	20	User5	16
User11	15	User11	45

Fig. 14

X	Y
10	5
20	16
15	45

Fig. 15

M1

Menu Options	Affinity
M2	0.266
M3	0.63
M4	0.2
M7	0.09
M11	0.5
M20	0.01

Fig. 16




Fig. 17

Menu	
M1	
M2	
M3	
<hr/>	
A1	
A2	
A3	

Fig. 18